Agent-based Modelling of Command-and-Control Effectiveness

Michael F. Ling

Theatre Operations Branch

Aeronautical and Maritime Research Laboratory, DSTO

Abstract

Uncertainty is one of the most pervasive elements in war, and it manifests itself no more prominently than in the area of Command and Control (C2). Superior information is no guarantee for its proper interpretation and utilisation. Likewise, well planned decisions at the HQ may well lose their value and relevance, or even be completely misinterpreted, on the way down the chain of command as a result of either the same information or order being viewed differently by different individuals, or the lower level commanders having very different perspectives because of their local circumstances, or both. The local commander may also be forced to make instant decisions based on gut reaction when the situation changes rapidly. In the dynamic environment of a military operation, uncertainty is highly difficult, if not impossible to control because of strong human interactions. New behaviours can emerge in these complex interactions. In this work, a high level agent-based model (ABM) is used to investigate how such uncertainties may hinder C2 operational effectiveness, even when high quality information is available. We also aim to understand some of the patterns that may emerge from the complex interactions between the agents.

Introduction

Over the past couple of decades there has been an increasing realisation [1,2] that conventional operations research tools based on rigorous mathematical equations and detailed physical description of combat cannot provide a realistic description of the complex and dynamic situations in which military operations are conducted. In such operations, may they be warfighting or peacekeeping, the participants have to interact with hostile or potentially hostile forces, and to respond to the hostile actions. In the process, a new situation or environment is created, which in turn triggers off new responses from both sides. Furthermore, war arouses some of the strongest human emotions which make it even more difficult to anticipate the behaviours of individuals in a command and control chain [3].

The main drawback of equation-based models (EBM) is that they are incapable of dealing with the dynamics of interactions between the combating sides and their reactions to each other's actions. Another serious challenge to EBM is that the world is fundamentally nonlinear, and consequently many problems defy the traditional scientific approach of analysis by decomposition. The nonlinearity of warfighting means that a small perturbation to some critical elements (initial conditions) can profoundly alter the outcomes, thus making reliable prediction extremely difficult, if not impossible. Finally, EBM cannot even begin to model the so-called intangibles such as human emotions, aggressiveness, fear, anger, team cohesion and so on [1].

With the advent of complexity theory and its application to warfare studies, it has become clear that warfare can be viewed as a complex adaptive system (CAS) which adapts, evolves and coevolves with its environment [2]. In this sense, a C2 system also behaves like a complex system which not only interacts with its environment but there are also dynamic (and often nonlinear) interactions between components of the system, namely the interactions between different levels of commanders and the commanders with the enemy. Conventional EBM are not equipped to deal with such ever changing conditions and nonlinearity. At best, one can use EBM to explore a range of possibilities, with a limited number of responses or contingency plans built into the wargaming. This, however, is far from being adequate for describing the behaviour of a system as it adapts and evolves with the changing environment.

Agent-based models (ABM) offer an opportunity to analyse the aforementioned complexity problems by concentrating on the behaviour of and interactions between the participating entities instead of the performance of specific weapons or sensors [2,4,5]. In other words, we shift our attention from analysing the performance of pieces of equipment to how different modes of operation may alter the outcome of a combat or peacekeeping or how the C2 system utilises information and acts upon it, and the consequence on the success or otherwise of a mission. In summary, we concentrate on the emergent pattern of the *whole* rather than the individual *parts*.

The use of intelligent agents to study the emergent behaviours of complex systems and operations, such as the roles played by training, aggressiveness or the effectiveness of command-and-control in combat performance, has attracted much attention in the military operations research (OR) community in recent years. Notable examples are models like the *Irreducible Semi-Autonomous Adaptive Combat* (ISAAC) [4,5] initiated by the US Marine Corps Combat Development Command (MCCDC) and the *Map Aware Non-uniform Automata* (MANA) [6,7] developed by the Defence Operational Technology Support Establishment (DOTSE), New Zealand.

The Model

This work makes use of the Map Aware Non-uniform Automata (MANA) developed in DOTSE. MANA was based heavily on Ilachinski's cellular automaton combat model ISAAC (Irreducible Semi-Autonomous Adaptive Combat) for the US Marine Corps Combat Development Command (MCCDC). The details of ISAAC has been discussed extensively by Ilachinski [4,5] and are available from the Center for Naval Analysis Website.

Briefly, ISAAC has approximately 30 parameters governing the behaviours and capabilities of the agents. They include weapon and sensor ranges and firepower, personalities such as each individual's inclination to move towards or away from its friend (cooperation) or enemy (aggressiveness), communication range between agents and squads. Since the moves in ISAAC are stochastic, a large number of simulations are needed in order to have a good statistical representation of the scenario outcomes.

MANA has fewer parameters than ISAAC [7] but also introduces two highly useful features. One is the trigger point which enables the analyst to model the change of behaviour under a wider range of circumstances than allowed in ISAAC. The other is the situational awareness (SA) map which plays the role of information centre for the combatants. The side with a greater situational awareness is aware of hostile threat over a greater distance, and is in a position to prepare itself for combat or different course of action earlier. It is in this capacity that the SA map will be used in this work for studying the relationship between C2 effectiveness, information superiority and combat performance.

The Baseline Scenario

The baseline scenario was set up with an intercept-combat missions in mind. A squad of Blue troops is given two tasks: to capture the Red Post within a given time period and to intercept any Red troops on route in order to prevent them from capturing a Blue Post (Red's next waypoint). These two tasks are given equal weight so that the Blues are not required to carry out an active search and destroy mission. They are only expected to intercept the Reds if they come within Blue's detection range (the Blues are supposed to have superior sensors and weapons). The MOE, measure of effectiveness, therefore takes into account the time taking to reach the Red Post, the number of Blue casualties and whether the Blue Post is taken by the Red agents.

Figure 1

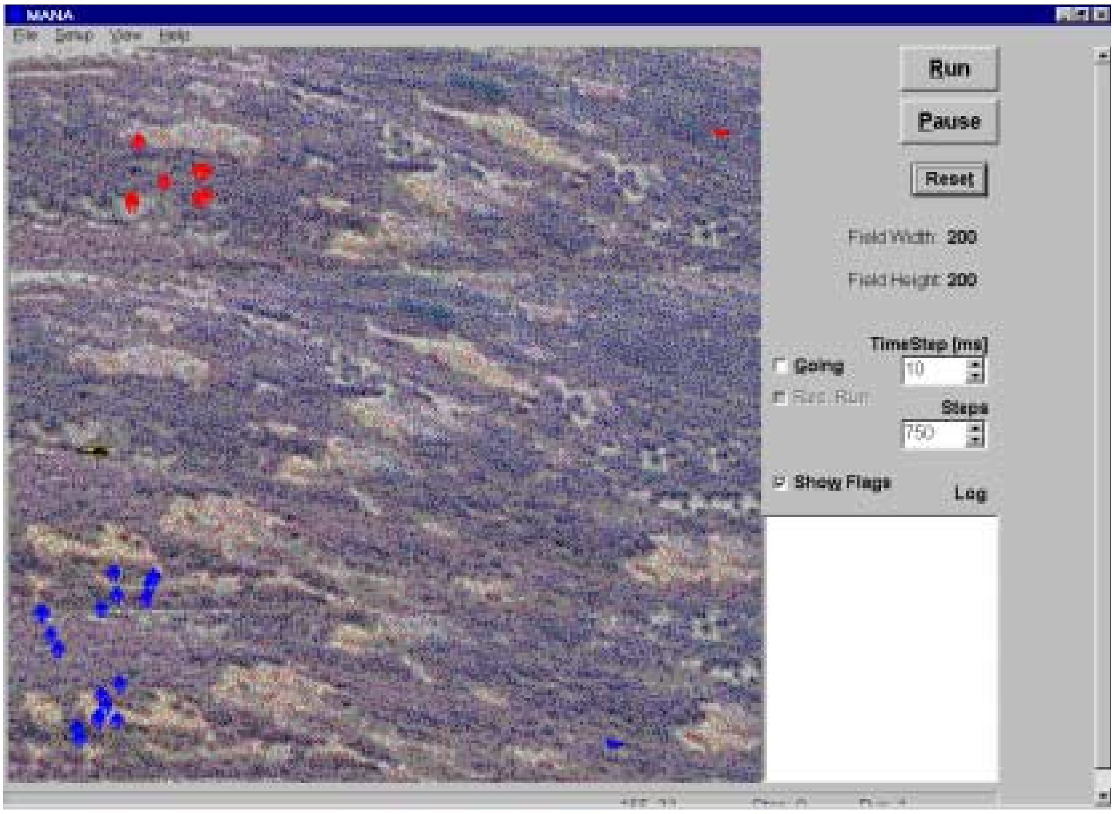

Figure 1 shows a typical MANA landscape with Blue and Red squads. The red and blue flags are the Red (Blue) and Blue (Red) Posts (waypoints), respectively. The helicopter symbol represents our UAV.

The baseline scenario begins with a Red Squad of 8 with relatively good training, firepower, fire range and sensor range, as well as good SA. On the blue side, there are 18 elite troops which are superior to the Red Squad in every respect. Furthermore, Blue has information superiority with the support of a UAV (unmanned aerial vehicle). In this scenario the SA value of the UAV is used as a reference point and taken be 100%.

The Experiment

In this study we are primarily concerned with two variables of the Blues: the value of situational awareness and the number of squads. The Blues' SA value can be viewed as their ability to interpret the superior information passed on from the UAV, whereas the number of squads on the Blue side represents the number of sub-commands. The former gives a measure of how well each Blue squad makes use of the information for decision-making in its local environment. The latter introduces extra randomness (uncertainty) to the way the Blues handle their tasks. It should be noted that all the actions on both Blue and Red sides represent only the functioning of a C2 system, and not necessarily military actions as such.

In MANA, each side has a Head Squad and all the other squads with the same allegiance (the same side) follow (obey) the Head Squad. For this study, the Head Squad becomes the highest command and all its subordinate squads play the role of sub-commanders with equal ranking. The idea is to observe the overall performance of the Blues as different SA values are allocated to the squads. While the current version of MANA restricts us to effectively only two levels of command, it is possible to emulate a C2 system using the squad structure, and to investigate the interactions between the Head Squad and its subordinate squads. It should be born in mind that every agent has a certain randomness (entropy) in its movement and decision-making process, thus making it possible to simulate the (unpredictable) actions and responses of human individuals with different perceptions of the same situation.

Throughout this study, the Red Squad has 8 members with an SA value of 67%. All the relevant parameters governing their performance, such as firepower, weapon range, sensor range, stealthiness are comparable but not as good as the corresponding values of the elite Blue Squad. In the case of the Blues, only the number of squads and their SA values will be altered.

In the first experiment an SA value of 86% is chosen for the Elite Blue Squad. It is thought unrealistic to expect even an elite squad to be able to make perfect use of information available from the UAV. The Blue Squad is sent from the lower left-hand corner of the MANA Map to capture the Red Post on the upper right-hand side. Almost diagonally cutting across their path, a Red Squad moves from the upper left-hand corner to the lower right-hand corner, making it almost certain that the Red will encounter the Blue on their way. See Figure 1. When the Blues traverses the Map, they encounter the Red and combat ensues. In each simulation run, if and when the Blues reach their destination, the number of casualties on both sides, and whether the Red Squad reaches its waypoint are noted and entered into our statistics.

In the second experiment, the Blue Squad SA is reduced to 60% and the same number of simulations and observations are repeated for this *non-elite* squad.

The next major step involves splitting the Blue Squad of 18 into three squads. The Head Squad has four members, one subordinate squad has 6 and the other has 8. We now have a clear command structure in which two squads are subordinate to the Head Squad, even though each squad has the freedom to make up its mind which task to carry out first due to the stochastic nature of the agents. In other words, splitting the Blues into three squads introduces greater complexity than the number may suggest.

First part of the this experiment with three Blue squads begin with all squads having the same SA value. While this new formation of squads consists of the same number of elite soldiers and same nominal combat capability and situational awareness, the way the tasks are carried out can be entirely different from that in the previous two experiments. For example, one squad may decide to move straight away to take over the Red Post while the other two choose to stay behind to intercept the Reds. The important question then is whether the sum of the parts is equal to the whole. In other words, are we likely to observe any emergent behaviour?

To complete our investigation into uncertainty in C2 and the relevance of information superiority in C2 and combat performance, we now give an SA value of 86% to the Head Squad but 50% and 33% to the other two Blue squads, thus simulating the effect of each commander interpreting information and orders differently, and quite possibly wrongly. Indeed, the randomness or entropy of each squad now truly plays the role of making its own decisions as each squad faces the combined tasks of moving to the next waypoint and intercepting the Reds.

Even though it is not realistic to expect the Blues to be able to make 100% use of available information from the UAV, two more experiments are conducted in which simulation runs are done on a single Ultra Elite Blue Squad and three Ultra Elite Blue Squads, all with 100% SA. The idea is both to use the results as reference points and to see if any emergent behaviours reveal themselves.

Results

The results of the six scenarios are shown in Table 1. Each set of scenario results are the average of forty simulations runs. While casualties are automatically recorded in a multiple simulation run in the current version of MANA, there is no distinction between the downing of the UAV and the killing a Blue squad member. Nor it is possible to automatically record the arrival of squads at their waypoint (capturing of Red or Blue Post). Therefore, all the simulation runs were watched carefully by the author on the computer screen as the agents moved and fought on the MANA landscape. While tedious and time-consuming, these direct observations had the added benefit of learning about the behaviours of the agents, thus helping to understand, for example, why attrition rate for the Blues is always higher in the 3-squad scenarios than the single-squad ones.

Note that the relatively probability of the UAV being killed should not be viewed with too much concern. The UAV is simply a symbolic representation of a source of superior information on the map of MANA. It has very high stealthiness but no firepower. Consequently, the Reds tend to shoot at the UAV repeatedly whenever the Blue agents are not shooting at them. Though it would have been a simple matter to give the UAV a greater tendency to move away from the Reds, the downing of the UAV was thought to be another test of the relevance of information superiority to the Blues' mission effectiveness.

Table 1

	Red killed (%)	Blue killed (%)	UAV killed (%)	RP taken (%)	BP taken (%)
1 Ultra Elite Blue Squad	99	0.4	17.5	83	0
1 Elite Blue Squad	89	0.6	40	80	0
1 Non-Elite Blue squad	66	0.7	45	100	33
3 Ultra Elite Blue Squads	96	1.7	15	100	5
3 Elite Blue Squads	95	1.4	14	100	2.3
3 Non-equal Blue Squads	63	2.7	15	100	2.5

Analysis

A convenient and effective way to analyse our results is by examining the columns of Table 1 which represent the measure of effectiveness (MOE). Recall that the MOE for the Blues includes capturing the Red Post within a fixed time period, intercepting the Red Squad and stopping them from reaching their waypoint (Blue Post), and avoiding casualties on the Blue side. While killing the Red agents is not explicitly part of the MOE, it has a clear connection with the SA value of the Blues.

It should be emphasised that the killing of the agents and the taking of posts are all merely measures of how well the C2 system carries out its tasks under the influence of uncertainty. It does not necessarily represent a real intercept-combat mission.

Even a cursory look down the columns in Table 1 shows that there is some rich information about the agent behaviours. Although the general trend is that the better the Blues' SA the better they perform, inconsistencies and surprises emerge all over the place. There are also some sharp variations in the attrition and capturing rates as one goes from 100% SA to 60% or lower. One may be tempted to dismiss some of the inconsistencies as statistically irrelevant. These figures, however, tell us only a small part of the story. The full picture of the agents' behaviours emerges only when they are monitored continuously on the computer screen. This time-consuming process is the reason that only a relatively small sample of 40 excursions were taken for each scenario.

In the case of a single Blue Squad, a higher SA value leads to higher kill rate, lower attrition loss and perfect record of preventing Blue Post from being taken. However, the Blues tend to spend too long engaging the enemy, with the result that nearly 20% of the time they could not reach their waypoint (Red Post) on time. With 100% SA, the Ultra Elite Blue Squad could usually locate the enemies very early and kill them straight away before moving on to their waypoint, leaving the Red agents little chance to attack the UAV, let alone capturing the Blue Post.

When the single-squad Blues are less alert to potential hostility, two interesting patterns emerge. First, the Red agents seem to become more cautious and tend to dig in and to try to shoot down the less threatening target of the UAV (it can't fire back) instead. Also, it is worth recalling that both sides now have comparable information of each other. Meanwhile, the Blues would take longer to locate the Reds and even longer to neutralise their threat. In some cases, the Blues would simply fail to notice the Reds completely and would move to take the Red Post right away, leaving the Reds ample opportunity to shoot down the UAV and even time to take the Blue Post as well. This behavioural pattern accounts for the sharp increase in UAV being downed and even sharper increase in Blue Post being captured as Blues' SA value drops by up to 40%, while at the same time the Red Post capturing rate goes up to 100%.

As already alluded to above, there is a great deal more complexity in the behaviour of the agents when the Blues are split into three squads. To begin with, the role of SA now becomes less transparent. First, the Blues now suffer much greater casualties than before. This is due to the fact that now only one or two Blue squads would engage the Reds, while at least one other squad would be free to take the Red Post. The result is that the Blues now have much smaller concentration of firepower and thus suffers much greater attrition loss. On the other hand, the success rate in reaching Red Post is now 100% regardless of their situational awareness. With the Blues now regularly getting killed and having enough firepower to eliminate the Reds, the Reds can reach their destination with greater regularity.

Two more patterns emerge in the 3-squad scenarios. First, since the three squads now always split up their tasks, there are at most two Blue squads at any time to engage the Reds. The Red agents tend to become more aggressive, and are now more inclined to shoot at the Blues as they don't have numerical superiority any longer, often resulting in a greater survival rate for the UAV because it has more time to move out of Red's weapon range.

Second, the 3-squad Ultra Elite Blues appear to be rather eager to engage the Reds regardless whether the numbers were in their favour or not. For example, the first squad that sees the enemy would engage it without waiting for the arrival and support of another squad. The consequence is that they suffer heavy loss without achieving the goal of preventing the Reds from capturing the Blue Post.

On a tactical note, the Blues need not kill the Reds to stop them from taking the Blue Post. It is sufficient for some of them to dig in and pose a constant threat to the Reds. The Red agents would not move until they feel no more threat from the Blues. By firing at the Reds the Blues exposed themselves and resulted in their own elimination. It does seem to confirm that foolish bravery does not pay.

Before closing, we should stress that models like MANA and ISAAC are stochastic in nature and therefore require a large number of simulations in order to get a reliable indication of the trend. It is also extremely informative to scan over the parameter space in order to study what is called the *fitness landscape*. For example, the MCCDC conducts their study on emergent behaviour in warfare, using ISAAC or similar tools, on a cluster of supercomputers in Maui, Hawaii, for generating large sample-parameter outputs with the hope of capturing nonlinearity and emergent properties. This is what MCCDC termed *data farming*.

Summary

It is clear that even a simple model like MANA can reveal some surprising patterns of behaviour of interacting agents in a simple C2 system. These interactions are dynamic and quite likely to be nonlinear, though the nonlinearity would be difficult to confirm without a much large sample space.

The main purpose of this work is to illustrate the potential application of high level agent-based models for identifying possible emergent properties in C2 operations. It shows that a C2 system is complex and dynamic. Even in a hugely simplified case, there are surprises and even nonlinear features emerging from the interactions among the components (*e.g.* different commanders) in a C2 system, and also the interactions between the system and its environment (*e.g.* the enemy). The so-called New Sciences of complex systems offer new methodologies and tools for investigating complex problems in military operations and systems, such as command and control, and may provide new insight into their workings.

This work may also contribute to understanding the role of information superiority in C2 effectiveness. First, any information is only as good as the people using it. The availability of high quality information is no good in the hands of poor commanders who either do not understand it or know how to use it properly. Moreover, it is clear that if superior information is not used intelligently, with appropriate adjustment in tactics and even doctrine, it may well be detrimental to mission effectiveness (see the results of 3-squad Ultra Elite Blue and the discussions).

Acknowledgment

Dr. Michael Lauren of DOTSE, New Zealand, kindly provided us with a copy of MANA which made this work possible in a very short space of time. There have been many useful discussions with Dr. Lauren. All these are gratefully acknowledged. The author thank DSTO for financial support to attend the 3rd Project Albert International Workshop in Auckland, New Zealand, in February 2001. This Workshop has proved immensely useful for getting this project off the ground.

References

1. Manoeuvre Warfare Science 2001, ed. Gary Horne and Mary Leonardi, Marine Corps Combat Development Command, Copyright © 2001.
2. Complexity, Global Politics, and National Security, ed. D.S. Alberts and T.J. Czerwinski, National Defense University, Washington. http://www.dodccrp.org.
3. John F. Schmitt, Command and (Out of) Control: The Military Implication of Complexity Theory, *ibid* Chapter 9.
4. Andrew Ilachinski, *Towards a Science of Experimental Complexity: An Artificial-Life Approach to Modelling Warfare*, http://www.can.org.isaac/.
5. Andrew Ilachinski, Irreducible Semi-Autonomous Adaptive Combat (ISAAC): An Artificial-Life Approach to Land Combat, Military Operations Research, **5**, No.3 29.
6. Michael K. Lauren, *Exploring the Value of Sensors to a Recce Unit Using Agent-based Models*, J. Battlefield Technology, 4, No.1, March 2001.
7. Michael K. Lauren and Richard T. Stephen, *Modelling Patrol Survivability in A Generic Peacekeeping Setting Using ISAAC*, DOTSE Report 177, NR 1358, ISSN 1174-3387.